1053.122

Disclosure ARC8-2000-0378

Prepared for and/or by an IBM Attorney - IBM Confidential

Created By: Dalit Naor **Created On:** 09/14/2000 01:36:11 PM
Last Modified By: Susana Delgado **Last Modified On:** 12/07/2000 01:21:06 PM

Required fields are marked with the asterisk (*) and must be filled in to complete the form .

***Title of disclosure (in English)**

A Subset-based Traitors Tracing Mechanism

Summary

Status	Under Evaluation
Processing Location	ARC
Functional Area	DPEB - Image & Multimedia Systems - (Robin Williams)
Attorney/Patent Professional	Marc D McSwain/Almaden/IBM
IDT Team	Robin Williams/Almaden/IBM; Hui-I Hsiao/Almaden/IBM; Bruce Lindsay/Almaden/IBM; C Mohan/Almaden/IBM; Spencer Ng/Almaden/IBM; Thomas Zimmerman/Almaden/IBM; Myron Flickner/Almaden/IBM; Anant Jhingran/Almaden/IBM; Marc D McSwain/Almaden/IBM
Submitted Date	11/16/2000 04:35:56 PM
Owning Division	RES
Lab	
Technology Code	
PVT Score	No PVT score has been calculated.To calculate a PVT score, press the 'Calculate' button.

Inventors with Lotus Notes IDs

Inventors: Dalit Naor/Almaden/IBM, Jeff Lotspiech/Almaden/IBM

Inventor Name	Inventor Serial	Div/Dept	Inventor Phone	Manager Name
> Naor, Dalit N.	3A7728	22/DPEM	457-1110	Lotspiech, Jeffrey B.
Lotspiech, Jeffrey B.	458235	22/DPEB	457-1851	Williams, Robin

> denotes primary contact

Inventors without Lotus Notes IDs

Simeon Naor
Serial Number : 210556 Company : IBM
Citizen of : IL
E-Mail : naor@wisdom.weizmann.ac.il
Business Address :
Almaden Research Center
650 Harry Road,
San Jose, CA. 95120
Business Phone : (408) 927-2737
Home Address :
247 Fulon Street
Palo Alto, CA. 94301

IDT Selection

IDT Team:	Attorney/Patent Professional:
Robin Williams/Almaden/IBM Hui-I Hsiao/Almaden/IBM Bruce Lindsay/Almaden/IBM C Mohan/Almaden/IBM Spencer Ng/Almaden/IBM Thomas Zimmerman/Almaden/IBM Myron Flickner/Almaden/IBM Anant Jhingran/Almaden/IBM Marc D McSwain/Almaden/IBM	Marc D McSwain/Almaden/IBM

Response Due to IP&L : 12/21/2000

***Main Idea**

1. Describe your invention, stating the problem solved (if appropriate), and indicating the advantages of using the invention.

This invention deals with prevention of piracy in the context of digital content distribution. Consider an encryption scheme whereby a Center broadcasts a message to a group of users so that only a subset of the users should be able to obtain the content of the message. Such schemes are naturally used for distribution of copyright protected content (such as music and movies) or for subscription-based systems (e.g. pay TV and Web Casting). A common problem with such schemes is that keys of certain users may leak and further be used by pirate decoders, software clones and other illegal means, thereby violating ownership rights of the data.

The invention provides a mechanism to combat the leakage of keys and their subsequent use by illegal decryption-boxes. Suppose that a pirate decryption-box contains the keys associated with at most t users $u_1, \ldots, u_t$ known as the "**traitors**". The goal of a tracing algorithm is to either

1. find the identities of those that contributed their keys to an illicit decryption box , or
2. render the box useless by finding a "pattern" that does not allow decryption using the box, but still allows broadcasting to the legitimate users.

When combined with an encryption scheme that is capable of revoking illegal users from future communications it yields a **trace-and-revoke** mechanism, which is a powerful tool to combat piracy. A tracing algorithm is evaluated based on (i) the number of illegal keys it is able to trace (ii) the level of performance downgrade it imposes on the encryption scheme (iii) the number of queries needed to trace the box.

The suggested scheme is a **black-box** tracing, i.e. one that does not take the decoder apart but by providing it with an encrypted message and observing its output (the decrypted message) tries to figure out who leaked the keys. It assumes that messages are encrypted using a **Subset-Cover** encryption scheme which satisfies the **bifurcation** property. The precise nature of such encryption schemes is defined below; two preferred embodiments for subset-cover revocation schemes having the bifurcation property are the Complete-Subtree method and the Subset-Difference method which are the subjects of Disclosure #....

Advantages of using this invention are:

1. In order to ~~trace t illegal users~~, it requires a message that consists of $t \log N$ keys where N is the total number of users in the entire system. A further improvement requires a message length of only $5t$

inventions. A Subset-Cover encryption scheme works as follows (as it covers all privileged users by smaller subsets).

- Each user u is initially assigned some secret information denote by I_u (typically, these are sets of keys).
- The scheme defines a collection of subsets of users $S_1, ..., S_w$ and their corresponding keys $K_1, ..., K_w$ so that for any $1 \leq i \leq w$ a user u can compute K_i from I_u if and only if it belongs to the subset S_i.
- Given M and P, The set P is partitioned into disjoint subsets $S = S\_i_1, S\_i_2, ..., S\_i_m$ so that every privileged user is in exactly one subset. M is then encrypted with the keys corresponding to these subsets:

$$< E_{K_{i1}}(K), E_{K_{i2}}(K), ..., E_{K_{im}}(K), F_K(M) >$$

This allows the users in P, and only them, to obtain M.

Bifurcation property

Our tracing mechanism requires that the Subset Cover algorithm satisfy the *bifurcation property*. The bifurcation property implies that for any subset S_i it is possible to partition S_i into two (or any constant) roughly equal sets and encrypt M using the two new sets instead of using S_i, i.e. there exist sets S_i_1 and S_i_2 such that

1. $S\_i = S\_i_1 \cup S\_i_2$
2. the size of S_i_1 is roughly the same as of S_i_2

For a Subset Cover scheme, let the **bifurcation value** be the relative size of the largest subset in such a split.

The two preferred embodiments for a Subset-Cover revocation scheme, the Complete Subtree and the Subtree Difference methods, satisfy the bifurcation property. In the case of the Complete Subtree Method the bifurcation value is 1/2 and for the Subtree Difference Method, the bifurcation value is 2/3.

Moreover, the Subtree Difference Method has an additional useful property: given any collection of r subsets $S\_i_1, ..., S\_i_r$, the method can cover all users that are *not* in $S\_i_1, ..., S\_i_r$ by at most 3r subsets.

In the discussion that follows the encryption scheme is viewed as a "box" that is capable of encrypting M when provided with either a specific partition S of all privileged users, or with the actual set P of privileged set of users. In the later, the partition that was used is also output. See diagram below.

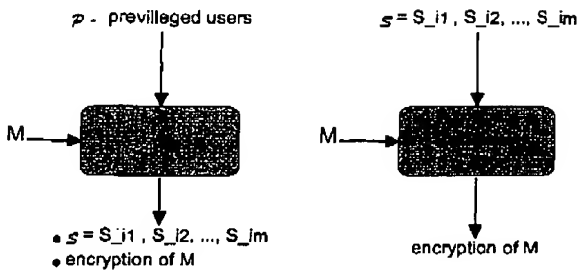

The Tracing Algorithm

Let *N* be the total number of users in the system. Suppose that a pirate decryption-box contains the keys associated with at most t users $u_1, \ldots, u_t$ known as the "**traitors**". The invention is a subsets-based **tracing algorithm**. It devises a sequence of queries that are given to the decoder whose result is either

- a subset of users consisting of the traitors, or
- a partition of users into subsets that renders the box useless, i.e. given a message that is encrypted with the given partition, the box decrypts the message with probability smaller than the threshold *q* while all good users can still decrypt.

Naturally, the tracing algorithm is based on constructing a useful sequence of partitions which will finally allow the detection of a traitor's identity.

An important procedure in our tracing mechanism is one that given a partition $S = S\_i_1, S\_i_2, \ldots, S\_i_m$ and an illegal box outputs one of two possible outputs: either

1. The box cannot decrypt when the encryption is done with partition S, or
2. Finds a subset S_i_j such that S_i_j contains a traitor.

Such a procedure is called **subset tracing.**

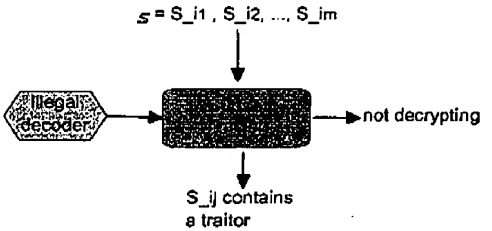

We explain our subset tracing procedure below. For now, let us assume that one exists, and we will now describe the general **tracing algorithm**, that uses the subset tracing procedure as a subroutine. The general algorithm maintains a partition $S = S\_i_1, S\_i_2, \ldots, S\_i_m$. At each phase one of the subsets is partitioned, and the goal is to partition a subset only if it contains a traitor. The initial partition is S = {all users}. A phase proceeds as follows:

At the beginning of the phase run the subset-tracing procedure with partition $S = S_{i_1}, S_{i_2}, ..., S_{i_m}$.

- If the procedure outputs that the box cannot decrypt with S then we are done, in the sense that we have found a way to disable the box without hurting any legitimate user.
- Otherwise,
  Let S_{i_j} be the set output by the subset-tracing procedure, namely S_{i_j} contains the a traitor.
  - If S_{i_j} contains only one possible candidate - it must be a traitor. Permanently revoke this user from the set of privileged users.
  - Otherwise, split S_{i_j} into two roughly equal subset and continue with the new partitioning. The existence of such a split is assured by the bifurcation property.

The number of iterations of the above can be at most $t \log_a N$, where a is the inverse of the bifurcation value.

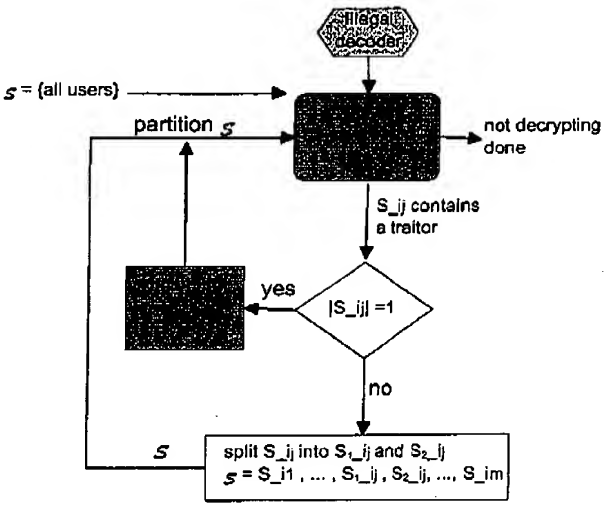

The Subset Tracing Procedure:

The Subset Tracing procedure first tests whether the box decodes a message that is legally encoded with the partition $S = S_{i_1}, S_{i_2}, ..., S_{i_m}$ with sufficient probability, say $p > 0.5$. By "legally encoded" we mean a normal message that would look exactly like normal operation. If the box does not decode, then it concludes (and outputs) that the box can not decrypt with S. Otherwise, it needs to find a subset S_{i_j} that contains a traitor.

Such a subset is found as follows. Let p_j be the probability that the box decodes the ciphertext

$$< E_{x_{i1}}(R_K), E_{x_{i2}}(R_K), ..., E_{x_{ij}}(R_K), E_{x_{ij+1}}(K), ..., E_{x_{im}}(K), F_K(M) >$$

where R_K is a random string of the same length as the key K; i.e., it is a false key. That is, p_j is the

probability of decoding when the first j subsets have false keys and the remaining subsets encode the correct key. If $| p_{j+1} - p_j | > p/m$ then it must be that S_j contains a traitor. We note that at least one such j always exists.

To efficiently find a subset that contains a traitor, employ the binary-search-like method described hereby that efficiently finds a pair of values p_j and p_{j+1} among $p_0, ..., p_m$ satisfying $| p_{j+1} - p_j | > p/m$. Starting with the entire interval [0,m], the search is repeatedly narrowed down to an arbitrary interval [a,b]. At each stage, the middle value $p_{a+b/2}$ is computed and the interval is further halved either to the left half or to the right half, depending on difference between $p_{a+b/2}$ and the endpoint values p_a and p_b of the interval. Observe that p_0 is p and p_1 is 0. Furthermore, in most practical cases, p is 1; in other words, the clone always decrypts during normal operation. The method is outlined below; it outputs the index j.

```
SubsetTracing(a,b,p_a,p_b)

    If (a=b-1)
        return b
    Else
        Let c = ⌈a+b/2⌉
        Compute p_c
        If | p_c - p_a | ≥ | (p_a - p_b) / 2 |
                SubsetTracing(a,c,p_a,p_c)
        Else
                SubsetTracing(c,b,p_c,p_b)
```

Efficiency: Subset tracing requires *O(log m)* evaluations of p_j. An evaluation of p_j must be within an accuracy that reveals a difference of the order of *1/m*; namely, p_j needs to be estimated so the difference between its true value and its estimated value does not exceed *1/2m* with assurance probability of 1-ε. Also , the true value of p_j can be as small as of the order of *1/m*. It follows from Chernoff bounds that $m^2 log (1/\varepsilon)$ ciphertext queries to the decoding box are sufficient to estimate such p_j within the required accuracy. Hence, a subset tracing procedure that works with success probability of ε log m requires $m^2 log m \ log (1/\varepsilon)$ ciphertext queries over the entire procedure.

Subset Tracing with Noisy Binary Search: It is possible to improve the efficiency of the subset tracing procedure by viewing it as a noisy-binary search procedure. The noisy binary search assumes that at each step of the decision tree the correct decision is obtained with probability 1-Q, where Q is a value close to 1/2, for example Q=1/3. In a model where each answer is correct with some fixed probability (say greater than 2/3) that is independent of history it is possible to perform binary search in $log N + log 1/Q$ queries where *log N* is the number levels in the search tree. Specifically for our case, can we assume that the computation of p_e at each step may yield a faulty value with probability Q. this yields that the number of the queries required over the entire procedure can be reduced to $m^2 (log m + log 1/Q)$.

Improving The Tracing Algorithm

Among the *t log N* subsets generated by the basic tracing algorithm, only *t* actually contain a traitor. The idea is to repeatedly merge those subsets which are not known to contain a traitor so as to reduce the number of subsets in the partition . For some encryption schemes it is possible to efficiently perform this merging, thus reducing the length of the message required to trace *t* traitors. For example, the preferred embodiment uses the Subset Difference method as the encryption scheme and requires a message of

only $5t$ to trace t traitors (instead of $t \log N$).

Specifically, we maintain at each iteration a **frontier** of at most 2t subsets and merge the rest of the subsets. In the following iteration a subset that contains a traitor is further partitioned; as a result, a new frontier is defined and the remaining subsets are re-grouped.

Frontier subsets

Let $S = S_{i_1}, S_{i_2}, \ldots, S_{i_m}$ be the partition at the current iteration. A pair of subsets $S_{i_{j1}}$ and $S_{i_{j2}}$ is said to be in the frontier if $S_{i_{j1}}$ and $S_{i_{j2}}$ resulted from a split-up of a single subset at an earlier iteration. Also neither $S_{i_{j1}}$ nor $S_{i_{j2}}$ was singled out by the subset tracing procedure so far. This definition implies that the frontier is composed of at most t disjoint pairs of **buddy subsets**.

The improved tracing algorithm proceeds in iterations. Every iteration starts with a partition $S = S_{i_1}, S_{i_2}, \ldots, S_{i_m}$. Denote by $F \subseteq S$ the frontier of S. An iteration consists of the following steps, by the end of which a new partition S' and a new frontier F' is defined.

- As before, use the Subset Tracing procedure to find a subset S_{i_j} that contains a traitor. If the tracing procedure outputs that the box can not decrypt with S then we are done. Otherwise, split S_{i_j} into S^1 and S^2.
- Set $F' = F \cup S^1 \cup S^2$ (include S^1 and S^2 in the new frontier). Furthermore, if S_{i_j} was in the frontier F and S_{i_k} was its buddy-subset in F then $F' = F' \setminus S_{i_k}$ (remove S_{i_k} from the new frontier).
- Compute a cover C for all receivers that are not covered by F'. Define the new partition S' as the union of C and F'.

An encryption method that can construct a small cover C for the non-frontier sets in the third step can take advantage of this improvement.

Tracing Traitors from Many Boxes

As new illegal decoding boxes, decoding clones and hacked keys are continuously being introduced during the lifetime of the system, a revocation strategy needs to be adopted in response. This revocation strategy is computed by first revoking the identities of all the receivers that need to be excluded, resulting in some partition S_0.

To trace traitors from possibly more than one illegal decoder and make all of these boxes non-decoding, the tracing algorithm needs to be run **in parallel** on all boxes by providing all boxes with the same input. The initial input is the partition S_0 that results from the set of all users that have not been revoked so far. As the algorithm proceeds, when the **first** box detects a traitor in one of the sets it re-partitions accordingly and the new partition is now input to **all** boxes simultaneously. The output of this simultaneous algorithm is a partition (or "revocation strategy") that renders all revoked receivers and illegal black boxes invalid.

3. If the same advantage or problem has been identified by others (inside/outside IBM), how have those others solved it and does your solution differ and why is it better?